

An informational approach to the global optimization of expensive-to-evaluate functions

Julien Villemonteix^{*†‡}, Emmanuel Vazquez[†] and Eric Walter^{*}

[†] *Supélec, Département Signaux et Systèmes Électroniques
91192 Gif-sur-Yvette, France*

^{*} *Laboratoire des Signaux et Systèmes, CNRS-Supélec-Univ Paris-Sud*

[‡] *Renault S.A., service 68240
78298 Guyancourt, France*

Abstract

In many global optimization problems motivated by engineering applications, the number of function evaluations is severely limited by time or cost. To ensure that each evaluation contributes to the localization of good candidates for the role of global minimizer, a sequential choice of evaluation points is usually carried out. In particular, when Kriging is used to interpolate past evaluations, the uncertainty associated with the lack of information on the function can be expressed and used to compute a number of criteria accounting for the interest of an additional evaluation at any given point. This paper introduces minimizer entropy as a new Kriging-based criterion for the sequential choice of points at which the function should be evaluated. Based on *stepwise uncertainty reduction*, it accounts for the informational gain on the minimizer expected from a new evaluation. The criterion is approximated using conditional simulations of the Gaussian process model behind Kriging, and then inserted into an algorithm similar in spirit to the *Efficient Global Optimization* (EGO) algorithm. An empirical comparison is carried out between our criterion and *expected improvement*, one of the reference criteria in the literature. Experimental results indicate major evaluation savings over EGO. Finally, the method, which we call IAGO (for Informational Approach to Global Optimization) is extended to robust optimization problems, where both the factors to be tuned and the function evaluations are corrupted by noise.

Key words: Gaussian process, global optimization, Kriging, robust optimization, stepwise uncertainty reduction

1 Introduction

This paper is devoted to global optimization in a context of expensive function evaluation. The objective is to find global minimizers in \mathbb{X} (the factor space, a bounded

subset of \mathbb{R}^d) of an unknown function $f : \mathbb{X} \rightarrow \mathbb{R}$, using a very limited number of function evaluations. Note that the global minimizer may not be unique (any global minimizer will be denoted as x^*). Such a problem is frequently encountered in the industrial world. For instance, in the automotive industry, optimal crash-related parameters are obtained using costly real tests and time-consuming computer simulations (a single simulation of crash-related deformations may take up to 24 hours on dedicated servers). It then becomes essential to favor optimization methods that use the dramatically scarce information as efficiently as possible.

To make up for the lack of knowledge on the function, surrogate (also called meta or approximate) models are used to obtain cheap approximations (Jones [2001]). They turn out to be convenient tools for visualizing the function behavior or suggesting the location of an additional point at which f should be evaluated in the search for x^* . Surrogate models based on Gaussian processes have received particular attention. Known in Geostatistics under the name of *Kriging* since the early 1960s (Matheron [1963]), Gaussian process models provide a probabilistic framework to account for the uncertainty stemming from the lack of information on the system. When dealing with an optimization problem, this framework allows the set of function evaluations to be chosen efficiently (Jones [2001], Jones et al. [1998], Huang et al. [2006]).

In this context, several strategies have been proposed, with significant advantages over traditional optimization methods when confronted to expensive-to-evaluate functions. Most of them *implicitly* seek a likely value for x^* , and then assume it to be a suitable location for a new evaluation of f . Yet, given existing evaluation results, the most likely location of a global minimizer is not necessarily a good evaluation point to improve our knowledge on x^* . As we shall show, by making full use of Kriging, it is instead possible to *explicitly* estimate the probability distribution of the optimum location, which allows an information-based search strategy.

Based on these observations, the present paper introduces minimizer entropy as a criterion for the choice of new evaluation points. This criterion, directly inspired from *stepwise uncertainty reduction* (Geman and Jedynak [1995]), is then inserted in an algorithm similar to the *Efficient Global Optimization* (EGO) algorithm (Jones et al. [1998]). We call the resulting algorithm IAGO, for *Informational Approach to Global Optimization*.

Section 2 recalls the principle of Kriging-based optimization, along with some general ideas on Gaussian process modeling that are used in Section 3 to build an estimate of the distribution of the global minimizers. Section 4 details the stepwise uncertainty reduction approach applied to global optimization, while Section 5 describes the corresponding algorithm and its extensions to noisy problems. Section 6 illustrates the behavior of the new algorithm on some simple benchmark problems, along with its performances compared with those of the classical EGO algorithm, chosen for its good compromise between local and global search (Sasena et al.

[2002]). Finally, after a conclusion section and to make this paper self-contained, Section 8 presents, as an appendix, some more results on Gaussian process modeling and Kriging.

2 Kriging-based global optimization

When dealing with expensive-to-evaluate functions, optimization methods based on probabilistic surrogate models (and Kriging in particular) have significant advantages over traditional optimization techniques, as they require fewer function evaluations. Kriging can indeed provide a cheap and accurate approximation of the function, but also an estimate of the potential error in this approximation. Numerous illustrations of this superiority can be found in the literature (see, for instance, Cox and John [1997]) and many variations have been explored (for extensive surveys, see Jones [2001] and Sasena et al. [2002]). As explained in this section, these methods deal with the cost of evaluation using an adaptive sampling strategy, replacing the optimization of the expensive-to-evaluate function f by a series of optimizations of a cheap criterion.

2.1 Gaussian process modeling and Kriging

This section briefly recalls the principle of Gaussian process (GP) modeling, and lays down the necessary notation. A more detailed presentation is available in the appendix (Section 8).

When modeling with Gaussian processes, the function f is assumed to be a sample path of a second-order Gaussian random process F . If we denote $(\Omega, \mathcal{A}, \mathcal{P})$ the underlying probability space, this amounts to assuming that $\exists \omega \in \Omega$, such that $F(\omega, \cdot) = f(\cdot)$. Whenever possible, we shall omit the dependence of F in ω to simplify notation.

In particular, given a set of n evaluation points $\mathbb{S} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ (the *design*), $\forall \mathbf{x}_i \in \mathbb{S}$ the evaluation result $f(\mathbf{x}_i)$ is viewed as a sample path of the random variable $F(\mathbf{x}_i)$. Kriging computes an unbiased linear predictor of $F(\mathbf{x})$ in the vector space $\mathbb{H}_{\mathbb{S}} = \text{span}\{F(\mathbf{x}_1), \dots, F(\mathbf{x}_n)\}$, which can be written as

$$\hat{F}(\mathbf{x}) = \boldsymbol{\lambda}(\mathbf{x})^\top \mathbf{F}_{\mathbb{S}}, \quad (1)$$

with $\mathbf{F}_{\mathbb{S}} = [F(\mathbf{x}_1), \dots, F(\mathbf{x}_n)]^\top$, and $\boldsymbol{\lambda}(\mathbf{x})$ the vector of Kriging coefficients for the prediction at \mathbf{x} .

Given the covariance of F , the Kriging coefficients can be computed along with the

variance of the prediction error

$$\hat{\sigma}^2(\mathbf{x}) = \text{var}(\hat{F}(\mathbf{x}) - F(\mathbf{x})). \quad (2)$$

The covariance of F is chosen within a class of parametrized covariances (for instance, the Matérn class), and its parameters are either estimated from the data or chosen a priori (see Section 8.3.2 for details on the choice of a covariance).

Once f has been evaluated at all evaluation points in \mathbb{S} , the predicted value of f at \mathbf{x} is given by

$$\hat{f}(\mathbf{x}) = \boldsymbol{\lambda}(\mathbf{x})^\top \mathbf{f}_{\mathbb{S}}, \quad (3)$$

with $\mathbf{f}_{\mathbb{S}} = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]^\top$ ($\mathbf{f}_{\mathbb{S}}$ is viewed as a sample value of $\mathbf{F}_{\mathbb{S}}$). The same results could be derived in a Bayesian framework, where $F(\mathbf{x})$ is Gaussian conditionally to the evaluations carried out ($\mathbf{F}_{\mathbb{S}} = \mathbf{f}_{\mathbb{S}}$), with mean $\hat{f}(\mathbf{x})$ and variance $\hat{\sigma}^2(\mathbf{x})$.

Note that

$$\forall \mathbf{x}_i \in \mathbb{S}, \quad \hat{F}(\mathbf{x}_i) = F(\mathbf{x}_i), \quad (4)$$

and that the prediction at $\mathbf{x}_i \in \mathbb{S}$ is $f(\mathbf{x}_i)$. When f is assumed to be evaluated exactly, Kriging is thus an interpolation, with the considerable advantage over other interpolation methods that it also provides an explicit characterization of the prediction error (zero-mean Gaussian with variance $\hat{\sigma}^2(\mathbf{x})$).

2.2 Adaptive sampling strategies

The general principle of optimization using Kriging is iteratively to evaluate f at a point that optimizes a criterion based on the model obtained using previous evaluation results. The simplest approach would be to choose a minimizer of the prediction \hat{f} as a new evaluation point. However, by doing so, too much confidence would be put in the current prediction and search is likely to stall on a local optimum (as illustrated by Figure 1). To compromise between local and global search, more emphasis has to be put on the prediction error that can indicate locations where additional evaluations are needed to improve confidence in the model. This approach has led to a number of criteria, based on both prediction and prediction error, designed to select additional evaluation points.

A standard example of such a criterion is *expected improvement* (EI). As the name suggests, it involves computing how much improvement in the optimum is expected, if f is evaluated at a given additional point. Let F be the Gaussian process model, as before, and f_{\min} be the current best function value obtained. The improvement expected from an additional evaluation of f at \mathbf{x} given $\mathbf{f}_{\mathbb{S}}$, the results of past evaluations, can then be expressed as

$$\text{EI}(\mathbf{x}) = \mathbb{E}[\max(f_{\min} - F(\mathbf{x}), 0) | \mathbf{F}_{\mathbb{S}} = \mathbf{f}_{\mathbb{S}}].$$




Fig. 1. Naive approach to optimization using Kriging: (*top*) prediction \hat{f} (bold line) of the true function f (dotted line, supposedly unknown) obtained from an initial design materialized by squares; (*bottom*) prediction after seven iterations of a direct minimization of \hat{f} .

Since $F(\mathbf{x})$ is conditionally Gaussian with mean $\hat{f}(\mathbf{x})$ and variance $\hat{\sigma}^2(\mathbf{x})$, a convenient expression appears (Schonlau [1997]):

$$\text{EI}(\mathbf{x}) = \hat{\sigma}(\mathbf{x}) \left[u\Phi(u) + \frac{d\Phi}{du}(u) \right], \quad (5)$$

with

$$u = \frac{f_{\min} - \hat{f}(\mathbf{x})}{\hat{\sigma}(\mathbf{x})}$$

and Φ the normal cumulative distribution function. The new evaluation point is then chosen as a global maximizer of $\text{EI}(\mathbf{x})$. An example is given on Figure 2, where the problem that deceived the naive method of Figure 1 is directly solved with the EI criterion. First introduced in (Jones et al. [1998]), this method has been used for computer experiments in (Sasena et al. [2002]), while modified criteria have been used in (Huang [2005]) and (Williams et al. [2000]) to deal with noisy functions.

In (Jones [2001]) and (Watson and Barnes [1995]), a fair number of alternative criteria are presented and compared. Although quite different in their formulation, they generally aim to answer the same question: What is the most likely position of \mathbf{x}^* ? Another, and probably more relevant, question is: Where should the evaluation be carried out optimally to improve our knowledge on the global minimizers?

In what follows, a criterion that addresses this question will be presented, along

with its performances. The reference for comparison will be EI, which is, according to Sasena et al. [2002], a reasonable compromise between local and global search, and has been successfully used in many applications.

3 Estimating the density of \mathbf{x}^*

Once a Kriging surrogate model \hat{f} has been obtained, any global minimizer of \hat{f} is a natural approximation of \mathbf{x}^* . However, it might be excessively daring to trust this approximation as it does not take in account the uncertainty of the prediction. A more cautious approach to estimating \mathbf{x}^* is to use the probabilistic framework associated with F . Of course, \mathbf{x}^* is not necessarily unique, and we shall focus on describing the set of all global minimizers of f as efficiently as possible.

3.1 Probabilistic modeling of the global minimizers of f

According to the GP model, a global minimizer \mathbf{x}^* of f corresponds to a global minimizer of this particular sample path of F . More formally, consider the *random* set $\mathcal{M}_{\mathbb{X}}^*$ of the global minimizers of F over \mathbb{X} , i.e. the set of all global minimizers




Fig. 2. EI approach to optimization using Kriging: (*top*) \hat{f} (bold line), 95% confidence intervals computed using $\hat{\sigma}$ (dashed line) and true function f (dotted line); (*bottom*) expected improvement.

for each sample path, which can be written as

$$\mathcal{M}_{\mathbb{X}}^*(\omega) = \{\mathbf{x}^* \in \mathbb{X} \mid F(\omega, \mathbf{x}^*) = \min_{\mathbf{u} \in \mathbb{X}} F(\omega, \mathbf{u})\}, \forall \omega \in \Omega,$$

with Ω the sample set. To ensure that $\mathcal{M}_{\mathbb{X}}^*$ is not empty, we assume that F has continuous sample paths with probability one. This continuity can be ensured through a proper choice of covariances (see, e.g., (Abrahamsen [1997])).

Let \mathbf{X}^* be a random vector uniformly distributed on $\mathcal{M}_{\mathbb{X}}^*$. The probability density function of this random vector conditional to past evaluation results, that we shall thereafter call conditional density of the global minimizers, is of great interest, as it allows one not only to estimate the global minimizers of f (for example, through the maximization of their probability density function), but also to characterize the uncertainty associated with this estimation. In fact, given the results $\mathbf{f}_{\mathbb{S}}$ of previous evaluations, the probability density function $p_{\mathbf{X}^*|\mathbf{f}_{\mathbb{S}}}(\mathbf{x})$ of \mathbf{X}^* conditionally to $\mathbf{f}_{\mathbb{S}}$ contains all of what has been assumed and learned about the system. However, no tractable analytical expression for such a quantity is available (Adler [2000], Sjö [2000]). To overcome this difficulty, the approach taken here is to consider a discrete version of the conditional distribution, and to approximate it using Monte Carlo simulations.

Let $\mathbb{G} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ be a finite subset of \mathbb{X} , $\mathcal{M}_{\mathbb{G}}^*$ be the random set of global minimizers of F over \mathbb{G} , and $\mathbf{X}_{\mathbb{G}}^*$ be a random vector uniformly distributed on $\mathcal{M}_{\mathbb{G}}^*$. The conditional probability mass function of $\mathbf{X}_{\mathbb{G}}^*$ given $\mathbf{f}_{\mathbb{S}}$ is then $\forall \mathbf{x} \in \mathbb{G}$

$$p_{\mathbf{X}_{\mathbb{G}}^*|\mathbf{f}_{\mathbb{S}}}(\mathbf{x}) = \mathbb{P}(\mathbf{X}_{\mathbb{G}}^* = \mathbf{x} \mid \mathbf{F}_{\mathbb{S}} = \mathbf{f}_{\mathbb{S}}).$$

It can be approximated using conditional simulations, *i.e.*, simulations of F that satisfy $\mathbf{F}_{\mathbb{S}} = \mathbf{f}_{\mathbb{S}}$. Assuming that non-conditional simulations are available, several methods exist to make them conditional (Chilès and Delfiner [1999]). Conditioning by Kriging seems the most promising of them in the present context and will be presented in the next section.

3.2 Conditioning by Kriging

This method, due to G. Matheron, uses the unbiasedness of the Kriging prediction to transform non-conditional simulations into simulations interpolating the results $\mathbf{f}_{\mathbb{S}}$ of the evaluations. Let Z be a zero-mean Gaussian process with covariance k (the same as that of F) and \hat{Z} be its Kriging predictor based on the random variables $Z(\mathbf{x}_i)$, $\mathbf{x}_i \in \mathbb{S}$, and consider the random process

$$T(\mathbf{x}) = \hat{f}(\mathbf{x}) + [Z(\mathbf{x}) - \hat{Z}(\mathbf{x})], \quad (6)$$

where \hat{f} is the mean of the Kriging predictor based on the design points in \mathbb{S} . Since this Kriging predictor is an interpolator, at evaluation points in \mathbb{S} , we have $\hat{f}(\mathbf{x}_i) = f(\mathbf{x}_i)$. Equation (4) implies that $Z(\mathbf{x}_i) = \hat{Z}(\mathbf{x}_i)$, which leads to $T(\mathbf{x}_i) = f(\mathbf{x}_i)$, $\forall \mathbf{x}_i \in \mathbb{S}$. In other words, T is such that all its sample paths interpolate the known values of f . It is then easy to check that T has the same finite-dimension distributions as F conditionally to past evaluation results (Delfiner [1977]). Note that the same vector $\boldsymbol{\lambda}(\mathbf{x})$ of Kriging coefficients is used to interpolate the data and the simulations at design points. Using (3), one can rewrite (6) as

$$T(\mathbf{x}) = Z(\mathbf{x}) + \boldsymbol{\lambda}(\mathbf{x})^\top [\mathbf{f}_{\mathbb{S}} - \mathbf{Z}_{\mathbb{S}}], \quad (7)$$

with $\mathbf{Z}_{\mathbb{S}} = [Z(\mathbf{x}_1), \dots, Z(\mathbf{x}_n)]^\top$.

In summary, to simulate F over \mathbb{G} conditionally to past evaluation results $\mathbf{f}_{\mathbb{S}}$, we can simulate a zero-mean Gaussian process Z over \mathbb{G} , and use the following method:

- Compute, for every point in \mathbb{G} , the vector of Kriging coefficients based on the design points in \mathbb{S} ,
- compute the Kriging prediction $\hat{f}(\mathbf{x})$ based on past evaluation results $\mathbf{f}_{\mathbb{S}}$ for every \mathbf{x} in \mathbb{G} ,
- sample over \mathbb{G} non-conditional sample paths of Z (provided that a Gaussian sampler is available, setting the proper covariance can be achieved using, for example, the Cholesky decomposition),
- apply (7) at every point in \mathbb{G} .

With this sampling method (see Figure 3 for an illustration), it becomes straightforward to estimate $p_{\mathbf{X}_{\mathbb{G}}^*|\mathbf{f}_{\mathbb{S}}}$. Let \mathbf{x}_i^* be a global minimizer of the i -th conditional simulation ($i = 1, \dots, r$) over \mathbb{G} (if it is not unique, choose one randomly). Then, for any \mathbf{x} in \mathbb{G} the empirical probability mass distribution

$$\hat{p}_{\mathbf{X}_{\mathbb{G}}^*|\mathbf{f}_{\mathbb{S}}}(\mathbf{x}) = \frac{1}{r} \sum_{i=1}^r \delta_{\mathbf{x}_i^*}(\mathbf{x}), \quad (8)$$

with δ the Kronecker symbol, tends almost surely towards $p_{\mathbf{X}_{\mathbb{G}}^*|\mathbf{f}_{\mathbb{S}}}(\mathbf{x})$ as r tends towards infinity. Moreover, $p_{\mathbf{X}_{\mathbb{G}}^*|\mathbf{f}_{\mathbb{S}}}$ tends in distribution towards $p_{\mathbf{X}^*|\mathbf{f}_{\mathbb{S}}}$ as \mathbb{G} becomes dense in \mathbb{X} .

Figure 4 presents the approximation achieved by $\hat{p}_{\mathbf{X}_{\mathbb{G}}^*|\mathbf{f}_{\mathbb{S}}}$ for an example where locating a global minimizer is not easy. Knowing the conditional distribution of $\mathbf{X}_{\mathbb{G}}^*$ gives valuable information on the areas of \mathbb{X} where a global minimizer might be located, and that ought to be investigated. This idea will be detailed in the next section.




Fig. 3. Conditioning a simulation: (*top*) unknown real curve f (dotted line), sample points (squares) and associated Kriging prediction \hat{f} (bold line); (*middle*) non-conditional simulation z , sample points and associated Kriging prediction \hat{z} (bold line); (*bottom*) the simulation of the Kriging error $z - \hat{z}$ is picked up from the non-conditional simulation and added to the Kriging prediction to get the conditional simulation (thin line).

4 The stepwise uncertainty reduction strategy

The knowledge about the global minimizers of f is summarized by the approximation $\hat{p}_{\mathbf{X}_{\mathbb{G}}^*|f_{\mathbb{S}}}$ of the conditional probability mass function of $\mathbf{X}_{\mathbb{G}}^*$. In order to evaluate the interest of a new evaluation of f at a given point, a measure of the expected information gain is required. An efficient measure is *conditional entropy*, as used in sequential testing (Geman and Jedynak [1995]) in the *Stepwise Uncertainty Reduction* (SUR) strategy. This section extends the SUR strategy to global optimization.

4.1 Conditional entropy

The entropy of a discrete random variable U (expressed in bits) is defined as:

$$H(U) = - \sum_u P(U = u) \log_2 P(U = u).$$

$H(U)$ measures the spread of the distribution of U . It decreases as this distribution gets more peaked. In particular :




Fig. 4. Estimation of the conditional point mass distribution of $\mathbf{X}_{\mathbb{G}}^*$: (*top*) Kriging interpolation, 95% confidence intervals and sample points; (*bottom*) estimated point mass distribution of $\mathbf{X}_{\mathbb{G}}^*$ using 10000 conditional simulations of F and a regular grid for \mathbb{G} .

- $\hat{p}_{\mathbf{X}_{\mathbb{G}}^* | f_{\mathbb{S}}}(\mathbf{x}) = 1/N \quad \forall \mathbf{x} \in \mathbb{G} \Rightarrow H(\mathbf{X}_{\mathbb{G}}^*) = \log_2(N),$
- $\hat{p}_{\mathbf{X}_{\mathbb{G}}^* | f_{\mathbb{S}}}(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} \neq \mathbf{x}_0 \\ 1 & \text{if } \mathbf{x} = \mathbf{x}_0 \end{cases} \Rightarrow H(\mathbf{X}_{\mathbb{G}}^*) = 0$

Similarly, for any event \mathcal{B} , the entropy of U relative to the probability measure $P(\cdot|\mathcal{B})$ is

$$H(U|\mathcal{B}) = - \sum_u P(U = u|\mathcal{B}) \log_2 P(U = u|\mathcal{B}).$$

The conditional entropy of U given another discrete random variable V is

$$H(U|V) = \sum_v P(V = v) H(U|V = v),$$

and the conditional entropy of U given \mathcal{B} and V is

$$H(U|\mathcal{B}, V) = \sum_v P(V = v|\mathcal{B}) H(U|\mathcal{B}, V = v). \quad (9)$$

More details on conditional entropy can be found in Cover and Thomas [1991].

4.2 Conditional minimizer entropy

Consider $F_Q(\mathbf{x})$ a discrete version of $F(\mathbf{x})$, defined as $F_Q(\mathbf{x}) = Q(F(\mathbf{x}))$ with Q a quantization operator. Q is characterized by a finite set of M real numbers $\{y_1, \dots, y_M\}$, and defined as

$$\forall u \in \mathbb{R} \quad Q(u) = \begin{cases} y_1 & \text{if } u \leq y_1 \\ y_i & \text{if } y_{i-1} < u \leq y_i \quad \forall i \in \llbracket 2, M \rrbracket \\ y_M & \text{if } y_M < u \end{cases}$$

For optimization problems, the SUR strategy for the selection of the next value of $\mathbf{x} \in \mathbb{X}$ at which f will be evaluated will be based on $H(\mathbf{X}_{\mathbb{G}}^* | \mathbf{F}_{\mathbb{S}} = \mathbf{f}_{\mathbb{S}}, F_Q(\mathbf{x}))$, the conditional entropy of $\mathbf{X}_{\mathbb{G}}^*$ given $F_Q(\mathbf{x})$ and the evaluation results $\{\mathbf{F}_{\mathbb{S}} = \mathbf{f}_{\mathbb{S}}\}$ (we shall refer to it later on as conditional entropy of the minimizer, or simply conditional entropy).

Using (9) we can write

$$H(\mathbf{X}_{\mathbb{G}}^* | \mathbf{F}_{\mathbb{S}} = \mathbf{f}_{\mathbb{S}}, F_Q(\mathbf{x})) = \sum_{i=1}^M P(F_Q(\mathbf{x}) = y_i | \mathbf{F}_{\mathbb{S}} = \mathbf{f}_{\mathbb{S}}) H(\mathbf{X}_{\mathbb{G}}^* | \mathbf{F}_{\mathbb{S}} = \mathbf{f}_{\mathbb{S}}, F_Q(\mathbf{x}) = y_i) \quad (10)$$

with

$$H(\mathbf{X}_{\mathbb{G}}^* | \mathbf{F}_{\mathbb{S}} = \mathbf{f}_{\mathbb{S}}, F_Q(\mathbf{x}) = y_i) = - \sum_{\mathbf{u} \in \mathbb{G}} p_{\mathbf{X}_{\mathbb{G}}^* | \mathbf{f}_{\mathbb{S}}, y_i}(\mathbf{u}) \log_2 p_{\mathbf{X}_{\mathbb{G}}^* | \mathbf{f}_{\mathbb{S}}, y_i}(\mathbf{u}) ,$$

and

$$p_{\mathbf{X}_{\mathbb{G}}^* | \mathbf{f}_{\mathbb{S}}, y_i}(\mathbf{u}) = P(\mathbf{X}^* = \mathbf{u} | \mathbf{F}_{\mathbb{S}} = \mathbf{f}_{\mathbb{S}}, F_Q(\mathbf{x}) = y_i).$$

$H(\mathbf{X}_{\mathbb{G}}^* | \mathbf{F}_{\mathbb{S}} = \mathbf{f}_{\mathbb{S}}, F_Q(\mathbf{x}))$ is a measure of the anticipated uncertainty remaining in $\mathbf{X}_{\mathbb{G}}^*$ given the candidate evaluation point \mathbf{x} and the result $\mathbf{f}_{\mathbb{S}}$ of the previous evaluations. Anticipation is introduced in (10) by considering the entropy of $\mathbf{X}_{\mathbb{G}}^*$ resulting from every possible sample value of $F_Q(\mathbf{x})$. At each stage of the iterative optimization, the SUR strategy retains for the next evaluation a point that minimizes the expected conditional minimizer entropy after the evaluation, *i.e.*, a point that maximizes the expected gain in information about $\mathbf{X}_{\mathbb{G}}^*$.

The conditional entropy of the minimizer thus takes in account the conditional statistical properties of F and particularly the covariance of the model. There lies the interest of the SUR strategy applied to global optimization. It makes use of what has been previously assumed and learned about f to pick up the most informative evaluation point. By contrast, the EI criterion (as most standard criteria) depends only on the conditional mean and variance of F at the design point considered.

5 Implementing the SUR strategy

5.1 IAGO algorithm

Our algorithm is similar in spirit to a particular strategy for Kriging-based optimization known as *Efficient Global Optimization* (EGO) (Jones et al. [1998]). EGO starts with a small initial design, estimates the parameters of the covariance of F and computes the Kriging model. Based on this model, an additional point is selected in the design space to be the location of the next evaluation of f using the EI criterion. The parameters of the covariance are then re-estimated, the model re-computed, and the process of choosing new points continues until the improvement expected from sampling additional points has become sufficiently small. The IAGO algorithm uses the same idea of iterative incorporation of the obtained information to the prior on the function, but with a different criterion.

The computation of the conditional entropy using (10) requires the choice of a quantization operator Q . We use the fact that $F(\mathbf{x})$ is conditionally Gaussian with known mean and variance (obtained by Kriging), to select a set of possible values $\{y_1, \dots, y_M\}$, such that

$$P(F_Q(\mathbf{x}) = y_i | \mathbf{F}_S = \mathbf{f}_S) = \frac{1}{M} \quad \forall i \in \llbracket 1 : M \rrbracket .$$

By doing so, we choose a different quantization operator Q for each value of \mathbf{x} to improve the precision with which the empirical mean of the entropy reduction over possible evaluation results is computed. This is simply an improved version of Monte Carlo integration. Here we used a set of ten possible values ($M = 10$).

For each of these possible values (or hypotheses $F(\mathbf{x}) = y_i$), $\hat{p}_{\mathbf{f}_S, y_i}$ is computed using conditional simulations. The conditional entropy is then obtained using (10). These operations are carried out on a discrete set of candidate evaluation points (see Section 5.2 for some details on the choice of this set), and a new evaluation of f is finally performed at a point that minimizes the conditional entropy. Next, as in the EGO algorithm, the covariance parameters are re-estimated and the model re-computed. The procedure for the choice of an additional evaluation point is described in Table 1.

When the number of additional function evaluations is not specified beforehand, we propose to use as a stopping criterion the conditional probability that the global minimum of the GP model be no further apart of the current minimum f_{\min}^n of the Kriging interpolation than a given tolerance threshold δ . The algorithm then stops when

$$P(F^* < f_{\min}^n + \delta | \mathbf{F}_S = \mathbf{f}_S) < P_{\text{Stop}} ,$$

Algorithm

Input: Set $\mathbb{S} = \{x_1, \dots, x_n\}$ of evaluation points and corresponding values $f_{\mathbb{S}}$ of the function f

Output: Additional evaluation point x_{new}

1. Choose \mathbb{G} , a discrete representation of \mathbb{X}
2. Set covariance parameters either a priori or by maximum-likelihood estimation based on $f_{\mathbb{S}}$
3. Compute r non-conditional simulations over \mathbb{G}
4. Compute $\hat{f}(x)$ and $\hat{\sigma}(x)$ over \mathbb{G} by Kriging from $f_{\mathbb{S}}$
5. **while** the set of candidate points has not been entirely explored
6. **do** Take an untried point x_c in the set of candidate points
7. Compute the parameters $\{y_1, \dots, y_M\}$ of the quantization operator Q
8. Compute $\Lambda = [\lambda(x_1), \dots, \lambda(x_N)]$ the matrix of Kriging coefficients at every point in \mathbb{G} based on evaluation points in \mathbb{S} and x_c
9. **for** $i \leftarrow 1$ **to** M
10. **do** Construct conditional simulations using (7) and assuming that $f(x_c) = y_i$
11. Find a global minimizer x_k^* of the k th conditional simulation over \mathbb{G} ($k = 1, \dots, r$)
12. Estimate $p_{X_{\mathbb{G}}^* | f_{\mathbb{S}}, y_i}$ over \mathbb{G} using (8)
13. Compute $H(X_{\mathbb{G}}^* | F_{\mathbb{S}} = f_{\mathbb{S}}, F_Q(x_c) = y_i)$
14. Compute the conditional entropy given an evaluation at x_c using (10)
15. **Output** x_{new} that minimizes the conditional entropy over the set of candidate points

Table 1

Selection of a new evaluation point for f .

with $F^* = \min_{x \in \mathbb{G}} F(x)$, and $P_{\text{Stop}} \in [0, 1]$ a critical value to be chosen by the user. Proposed by Schonlau (Schonlau [1997]), this stopping criterion is well suited here, since evaluating the repartition function of $f(x^*)$ does not require any additional computation.

5.2 Computational complexity

With the previous notation, n the number of evaluation points, r the number of conditional simulations, N the number of points in \mathbb{G} and M number of possible results for an evaluation, the computational complexity for the approximation of conditional entropy (Steps 7 to 14 in Table 1) is as follows:

- computing Kriging coefficients at every point in \mathbb{G} (Step 8): $O(nN)$ (once the covariance matrix in (17) has been factorized, Kriging at an untried point is simply in $O(n)$),
- constructing conditional simulations (Step 10): $O(nrN)$ (M is actually not involved since the main part of conditioning can be carried out outside the loop on the possible evaluation values),
- locating the global minimizers for each simulation by exhaustive search (Step 11):

$$O(rNM),$$

Since all other operations are in $O(N)$ at most, evaluating conditional entropy at any given point requires $O(N)$ operations.

To complete the description of an implementable algorithm, we must specify a choice for \mathbb{G} and a policy for the minimization of conditional entropy. What follows is just an example of a possible strategy, and many variants could be considered.

The simplest choice for \mathbb{G} is a uniform grid on \mathbb{X} . However, as the number of evaluations of f increases, the spread of $p_{\mathbf{x}_{\mathbb{G}}^*|f_{\mathbb{S}}}$ diminishes along with the precision for the computation of the entropy. To keep a satisfactory precision over time, \mathbb{G} can be a random sample of points in \mathbb{X} , re-sampled after every evaluation of f with the density $\hat{p}_{\mathbf{x}_{\mathbb{G}}^*|f_{\mathbb{S}}}$. Re-sampling makes it possible to use a set \mathbb{G} with a smaller cardinal and to escape, at least partly, the curse of dimensionality (to resample using $\hat{p}_{\mathbf{x}_{\mathbb{G}}^*|f_{\mathbb{S}}}$, any non-parametric density estimator could be used along with a sampling method such as Metropolis Hastings, see, e.g., (Chib and Greenberg [1995])).

Ideally, to choose an additional evaluation point for f using IAGO, conditional entropy should be minimized over \mathbb{X} . However, this of course is in itself a global optimization problem, with many local optima. It would be possible to design an ad-hoc optimization method (as in (Jones [2001])), but this perspective is not explored here. Instead, we evaluate the criterion extensively over a chosen set of candidate points. Note that only the surrogate model is involved at this stage, which makes the approach practical. The idea is, exactly as for the choice of \mathbb{G} , to use a space-filling sample covering \mathbb{X} and resampled after each new evaluation. The current implementation of IAGO simply uses a Latin Hyper Cube (LHC) sample, however, it would be easy to adapt this sample iteratively using the conditional distribution of the minimizers $\hat{p}_{\mathbf{x}_{\mathbb{G}}^*|f_{\mathbb{S}}}$ as a prior. For instance, areas of the design space where the density is sufficiently small could be ignored. After a few evaluations, a large portion of the design space usually satisfies this property, and the computations saved could be used to improve knowledge on the criterion by sampling where $\hat{p}_{\mathbf{x}_{\mathbb{G}}^*|f_{\mathbb{S}}}$ is high (using the same approach as for the choice of \mathbb{G}).

As dimension increases, trying to cover the factor space while keeping the same accuracy leads to an exponential increase in complexity. However, in a context of expensive function evaluation, the objective is less to specify exactly all global minimizers (which could be too demanding in function evaluations anyway), than to use available information efficiently to reduce the likely areas for the location of these minimizers. This is exactly the driving concept behind IAGO. In practice, within a set of one thousand candidate points, picking an additional evaluation point requires about five minutes with a standard personal computer (and this figure is relatively independent of the dimension of the factor space). Moreover, the result obtained can be trusted to be a consistent choice within this set of candidate points, in regard of what has been assumed and learned about f .

5.3 Taking noise in account

Practical optimization problems often involve noise. This section discusses possible adaptations of the optimization algorithm that make it possible to deal with noisy situations, namely noise on the evaluation of f and noise on the factors.

5.3.1 Noise on the evaluation of f

When the evaluations of f are corrupted by noise, the algorithm must take this fact into account. A useful tool to deal with such situations is *non-interpolative Kriging* (see Section 8.2).

If the evaluation at $\mathbf{x}_i \in \mathbb{S}$ is assumed to be corrupted by an additive Gaussian noise ε_i with known mean and variance, the Kriging prediction should no longer be interpolative (see Section 8.2). The optimization algorithm remains nearly unchanged, except for the conditional simulations. We now wish to build sample paths of F conditionally to evaluation results, i.e. realizations of the random variables $f(\mathbf{x}_i) + \varepsilon_i$ for $\mathbf{x}_i \in \mathbb{S}$. Since the variance of the prediction error is no longer zero at evaluation points (in other words, there is some uncertainty left on the values of f at evaluation points), we first have to sample, at each evaluation points, from the distribution of F conditionally to noisy evaluation results. An interpolative simulation, based on these sample values, is then built using conditioning by Kriging. An example of such a simulation is proposed on Figure 5.

5.3.2 Noise on the factors

In many industrial design problems, the variability of the values of the factors in mass production has a significant impact on the performance that can be achieved. In such a case, one might want to design a system that optimizes some performance measure while ensuring that the performance uncertainty (stemming from noise on the factors) remains under control. These so-called *robust optimization* problems can generally be written as

$$\arg \min_{\mathbf{x} \in \mathbb{D}} J(\mathbf{x}) , \quad (11)$$

with $J(\mathbf{x})$ a cost function reflecting some statistical property of the corrupted performance measure $f(\mathbf{x} + \varepsilon)$, where ε is a random vector accounting for noise on the factors. Classical cost functions are:

- mean: $J(\mathbf{x}) = \mathbb{E}_{\varepsilon}[f(\mathbf{x} + \varepsilon)]$,
- standard deviation: $J(\mathbf{x}) = \sigma_{\varepsilon}[f(\mathbf{x} + \varepsilon)] = \sqrt{\text{var}(f(\mathbf{x} + \varepsilon))}$,
- linear combination of mean and standard deviation: $J(\mathbf{x}) = \mathbb{E}_{\varepsilon}[f(\mathbf{x} + \varepsilon)] + k\sigma_{\varepsilon}[f(\mathbf{x} + \varepsilon)]$,




Fig. 5. Example of prediction by Kriging (bold line) of noisy measurements represented by squares. Dashed lines represent 95% confidence regions for the prediction and the thin solid line is an example of conditional simulation obtained using the method presented in Section 5.3.1 (crosses represent the simulated measurements.)

- α -quantile: $J(\mathbf{x}) = Q^\alpha(\mathbf{x})$
with $Q^\alpha(\mathbf{x})$ such that $P(f(\mathbf{x} + \boldsymbol{\varepsilon}) < Q^\alpha(\mathbf{x})) = \alpha$.

Using, for example, the α -quantile as a cost function, it is possible to adapt our optimization algorithm to solve (11). Given a set of evaluation results \mathbf{f}_S at noise-free evaluation points, and assuming that it is possible to sample from the distribution p_ε of ε , a Monte-Carlo approximation $\hat{Q}^\alpha(\mathbf{x})$ of $Q^\alpha(\mathbf{x})$ is easily obtained by computing $\hat{f}(\mathbf{x} + \varepsilon)$ over a set sampled from p_ε . The global optimization algorithm can then be applied to $Q^\alpha(\mathbf{x})$ instead of f , using pseudo-evaluations $\hat{\mathbf{Q}}_S^\alpha = [\hat{Q}^\alpha(\mathbf{x}_1), \dots, \hat{Q}^\alpha(\mathbf{x}_n)]$ instead of \mathbf{f}_S .

It is of course possible to combine these ideas and to deal simultaneously with noise both on the factors and the function evaluations.

6 Illustrations

This section contains some simple examples of global optimization using IAGO, with a regular grid as a set of candidate evaluation points. An empirical comparison with global optimization using expected improvement is also presented. The

Matérn covariance class will be used for Kriging prediction, as it facilitates the tuning of the variance, regularity and range of correlation of the underlying random process, but note that any kind of admissible covariance could have been used. The parameters of the covariance may be estimated from the data using a maximum-likelihood approach (see Section 8.3).

6.1 A one-dimensional example


Consider the function with two global minimizers illustrated by Figure 6 and defined by $f : x \mapsto 4[1 - \sin(x + 8 \exp(x - 7))]$. Given an initial design consisting of three points, the IAGO algorithm is used to compute six additional points iteratively. The final Kriging model is depicted in the left part of Figure 6, along with the resulting point mass conditional distribution for the minimizer on the right part. After adding some noise on the function evaluations, the variant of the algorithm presented in Section 5.3.1 is also applied to the function with the same initial design. In both cases, the six additional evaluations have significantly reduced the uncertainty associated with the position of the global minimizers. The remaining likely locations reduce to two small areas centered on the two actual global minimizers. In the noisy case, larger zones are identified, a direct consequence of the uncertainty associated with the evaluations.

Figure 7 illustrates robust optimization using the same function and initial design. The cost function used is the 90%-quantile $Q^{90\%}$, which is computed on the surrogate model but also, and only for the sake of comparison, on the true function using Monte Carlo uncertainty propagation (the quantile is approximated using 5000 simulations). After six iterations of the robust optimization algorithm, the distribution of the robust minimizer is sufficiently peaked to give a good approximation of the true global robust minimizer.


These results are encouraging as they show that the requirement of a fast uncertainty reduction is met. The next section provides some more examples, along with a comparison with EGO, the EI-based global optimization algorithm.

6.2 Empirical comparison with expected improvement


Consider first the function described by Figure 8. Given an initial design of three points, both the EI and conditional entropy are computed. Their optimization provides two candidate evaluation points for f , which are also presented on Figure 8, along with the post-evaluation prediction and conditional point mass distribution for $\mathbf{X}_{\mathbb{G}}^*$. For this example, the regularity parameter of the Matérn covariance is set a priori to a high value (2.5), and it is therefore likely that two evaluations close to one another, as proposed by the IAGO algorithm, will give some valuable informa-



(a) Kriging prediction and point mass distribution of the global minimizers based on the initial design



(b) Standard IAGO algorithm (noise free case)



(c) IAGO algorithm for noisy evaluations (the additive noise is zero-mean Gaussian with standard deviation 0.2)

Fig. 6. Example of global optimization using IAGO on a function of one variable (dotted line), with an initial design consisting of three points (represented by squares). Six additional evaluations are carried out (triangles) using two versions of the IAGO algorithm. The graphs on the *left part* of the figure account for final predictions, while the *right part* presents the final point mass distributions of the global minimizers

tion about the part of \mathbb{X} located on the left of the evaluation point, while improving the characterization of the global minimizer. By taking in account the covariance of F through conditional simulations, the conditional entropy uses regularity to conclude faster. The resulting conditional point mass distribution of the minimizers is then generally more peaked using the IAGO algorithm than using the EGO





Fig. 7. Example of robust optimization using IAGO and the cost function $Q^{90\%}$. The function f (dotted line), corrupted by a Gaussian noise on the factor (zero mean with a standard deviation of 0.2), is studied starting from an initial design of three points (as in Figure 6). Six additional evaluations are carried out (triangles), which are used to estimate the cost function based on the Kriging model (bold line), along with the estimated point mass distribution of the robust minimizers (*right*). The cost function $Q^{90\%}$ estimated, only for the sake of comparison, from the true function using Monte Carlo uncertainty propagation is also provided (mixed line).

algorithm (as illustrated by Figure 8(c) and Figure 8(b)).


Consider now the Branin function (see, for instance, (Dixon and Szegö [1978])), defined as

$$\begin{aligned}
 f : [-5, 10] \times [0, 15] &\longrightarrow \mathbb{R} \\
 (x_1, x_2) &\longmapsto \left(x_2 - \frac{5.1}{4\pi^2} x_1^2 + \frac{5}{\pi} x_1 - 6 \right)^2 \\
 &\quad + 10 \left(1 - \frac{1}{8\pi} \right) \cos(x_1) + 10 .
 \end{aligned}$$


It has three global minimizers $\mathbf{x}_1^* \approx (-3.14, 12.27)^\top$, $\mathbf{x}_2^* \approx (3.14, 2.27)^\top$ and $\mathbf{x}_3^* \approx (9.42, 2.47)^\top$, and the global minimum is approximately equal to 0.4. Given an initial uniform design of fifteen points, fifteen additional points are iteratively selected and evaluated using the IAGO and EGO algorithms. These parameters are estimated on the initial design, and kept unchanged during both procedures. The positions of these points are presented on Figure 9 (left), along with the three global minimizers. Table 2 summarizes the results obtained with EGO and IAGO, based on the final Kriging models obtained with both approaches. Note that the EI criterion in EGO is maximized with a high precision, while the conditional entropy in IAGO is computed over a thousand candidate evaluation points located on a regular grid. It appears nevertheless that the algorithm using EI stalls on a single global minimizer, while the conditional entropy allows a relatively fast estimation of all three of them. Besides, the search is more global with IAGO, which yields a better global approximation of the supposedly unknown function. If twenty additional evaluations are carried out (as presented in the right part of Figure 9), the final Kriging prediction using the SUR approach estimates the minimum with an error of less than 0.05 for all three minimizers (cf. Table 2), while the EI criterion does



(a) Initial prediction and density of the global minimizers



(b) Prediction and density of the global minimizers after an additional evaluation of f chosen with EI



(c) Prediction and density of the global minimizers after an additional evaluation of f chosen with conditional entropy

Fig. 8. Comparison between conditional entropy and EI: the *left* side contains the Kriging predictions before and after an additional evaluation chosen with either EI or conditional entropy, while the *right* side presents the corresponding conditional density of the global minimizers.

not improve the information on any minimizer any further. The difference between the two strategies is clearly evidenced. The EI criterion, overestimating the confidence in the initial prediction, has led to performing evaluations extremely close to one another, for a very small information gain. In a context of expensive function evaluation, this is highly detrimental. The entropy criterion, using the same covariance parameters, does not stack points almost at the same location before having identified the most likely zones for the minimizers. The use of what has

	EGO		IAGO	
	15 iterations	35 iterations	15 iterations	35 iterations
Euclidean distance between \mathbf{x}_1^* and its final estimate	3.22	3.22	2.18	0.23
Value of the true function at estimated minimizer	17.95	17.95	2.59	0.40
Euclidean distance between \mathbf{x}_2^* and its final estimate	2.40	2.40	0.44	0.18
Value of the true function at estimated minimizer	13.00	13.00	0.85	0.42
Euclidean distance between \mathbf{x}_3^* and its final estimate	0.04	0.04	0.82	0.23
Value of the true function at estimated minimizer	0.40	0.40	1.94	0.44

Table 2

Estimation results for the Branin function using evaluations of Figure 9

been assumed and learned about the function is clearly more efficient in this case, and this property should be highly attractive when dealing with problems of higher dimension.

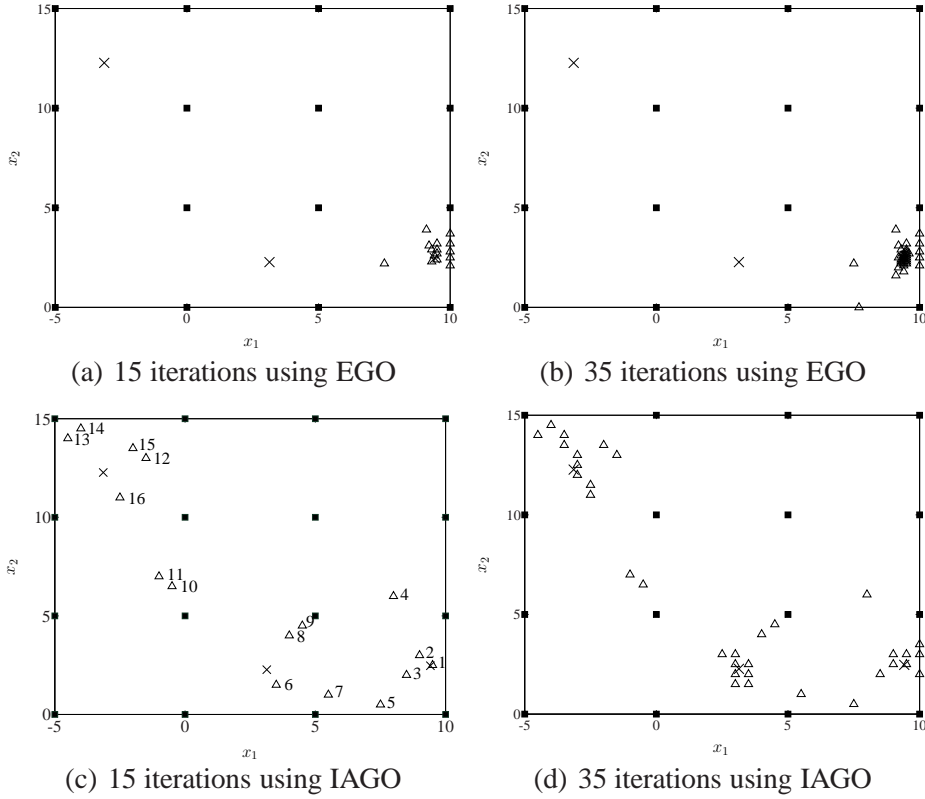


Fig. 9. fifteen iterations of two optimization algorithms, that differ by their criteria for selecting evaluation points for f , on the Branin function: (*top*) the EI criterion is used, (*bottom*) the conditional entropy criterion is used with a thousand candidate evaluation points for f set on a regular grid (squares account for initial data, triangles for new evaluations, and crosses give the actual locations of the three global minimizers).

7 Discussion

7.1 Robustness to uncertainty on the covariance parameters

Jones [2001] studies the potential of Kriging-based global optimization methods such as EGO. One of his most important conclusions, is that these methods “*can perform poorly if the initial sample is highly deceptive*”. An eloquent example is provided on page 373, where a sinus function is sampled using its own period, leading to a flat prediction over the domain, associated with a small prediction error.

This potential for deception is present throughout the IAGO procedure, and should not be ignored. To overcome this difficulty, several methods have been proposed (see, e.g., the Enhanced Method 4 in (Jones [2001]) or (Gutmann [2001])), which achieve some sort of robustness to an underestimation of the prediction error and more generally to a bad choice of covariance. They seem to perform better than classical algorithms, including EGO.

Comparing the IAGO approach to such methods is an interesting topic for future research. The issue considered here was to demonstrate the interest of the conditional entropy criterion, and we feel that this should be done independently from the rest of the procedure.

It is of course essential to make the IAGO algorithm robust to errors in the estimation of the covariance parameters. In many industrial problems, this can be easily done by using prior knowledge on the unknown function to restrict the possible values for these parameters. For example, experts of the field often have information regarding the range of values attainable by the unknown function. This information can be directly used to restrict the search space for the variance of the modeling process F , or even to choose it beforehand.

More generally, given the probabilistic framework used here, it should be relatively easy to develop a Bayesian or minimax extension of the IAGO algorithm to guide the estimation of the parameters of the covariance. A comparison with robust methods such as those detailed in (Jones [2001]) will then be essential.

7.2 Conclusions and perspectives

In this paper, a stepwise uncertainty reduction strategy has been used for the sequential global optimization of expensive-to-evaluate functions. This strategy iteratively selects a minimizer of the conditional minimizer entropy as the new evaluation point \mathbf{x}_{new} . To compute this entropy, a Gaussian random model of the function

evaluations is used and the minimizer entropy is estimated through Kriging and conditional simulations. At each iteration, the result of the evaluation at \mathbf{x}_{new} is incorporated in the data base used to re-build the Kriging model (with a possible re-estimation of the parameters of its covariance).

We have shown on some simple examples that, compared to the classical EI-based algorithm EGO, the method proposed significantly reduces the evaluation effort in the search for global optimizers. The SUR strategy allows the optimization method to adapt the type of search to the information available on the function. In particular, the conditional entropy criterion makes full use of the assumed regularity of the unknown function to balance global and local searches.

Choosing an adequate set of candidate points is a crucial point, as it must allow a good estimation of a global minimizer of the criterion, while keeping computation feasible. Promising results have already been obtained with space-filling designs, and adaptive sampling based on the conditional density of the global minimizers should provide useful results as dimension increases.

Extension to constrained optimization is an obviously important topic of future investigations. When it is easy to discard the candidate points in \mathbb{X} that do not satisfy the constraints, the extension is trivial. For expensive-to-evaluate constraints, the extension is a major challenge.

Finally, the SUR strategy associated with conditioning by Kriging is a promising solution for the robust optimization of expensive-to-evaluate functions, a problem that is central to many industrial situations, for which an efficient product design must be found in the presence of significant uncertainty on the values actually taken by some factors in mass production. In addition, robustness to the uncertainty associated with the estimation of the parameters of the covariance should also be sought.

8 Appendix: modeling with Gaussian processes

This section recalls the main concepts used in this paper, namely Gaussian process modeling and Kriging. The major results will be presented along with the general framework for the estimation of the model parameters.

8.1 Kriging when f is evaluated exactly

Kriging (Matheron [1963], Chilès and Delfiner [1999]) is a prediction method based on random processes that can be used to approximate or interpolate data. It can also be understood as a kernel regression method, such as *splines* (Wahba

[1998]) or *Support Vector Regression* (Smola [1998]). It originates from geostatistics and is widely used in this domain since the 60s. Kriging is also known as the *Best Linear Unbiased Prediction* (BLUP) in statistics, and has been more recently designated as Gaussian Processes (GP) in the 90s in the machine learning community.

As mentioned in Section 2.1, it is assumed that the function f is a sample path of a second-order Gaussian random process F . Denote by $m(\mathbf{x}) = E[F(\mathbf{x})]$ the mean of $F(\mathbf{x})$ and by $k(\mathbf{x}, \mathbf{y})$ its covariance, written as

$$k(\mathbf{x}, \mathbf{y}) = \text{cov}(F(\mathbf{x}), F(\mathbf{y})).$$

Kriging then computes the BLUP of $F(\mathbf{x})$, denoted by $\hat{F}(\mathbf{x})$, in the vector space generated by the evaluations $\mathbb{H}_{\mathbb{S}} = \text{span}\{F(\mathbf{x}_1), \dots, F(\mathbf{x}_n)\}$. As an element of $\mathbb{H}_{\mathbb{S}}$, $\hat{F}(\mathbf{x})$ can be written as

$$\hat{F}(\mathbf{x}) = \boldsymbol{\lambda}(\mathbf{x})^\top \mathbf{F}_{\mathbb{S}}, \quad (12)$$

As the BLUP, $\hat{F}(\mathbf{x})$ must have the smallest variance for the prediction error

$$\hat{\sigma}^2(\mathbf{x}) = \mathbb{E}[(\hat{F}(\mathbf{x}) - F(\mathbf{x}))^2], \quad (13)$$

among all unbiased predictors. The variance of the prediction error satisfies

$$\hat{\sigma}^2(\mathbf{x}) = k(\mathbf{x}, \mathbf{x}) + \boldsymbol{\lambda}(\mathbf{x})^\top \mathbf{K} \boldsymbol{\lambda}(\mathbf{x}) - 2\boldsymbol{\lambda}(\mathbf{x})^\top \mathbf{k}(\mathbf{x}), \quad (14)$$

with

$$\mathbf{K} = (k(\mathbf{x}_i, \mathbf{x}_j)), \quad (i, j) \in \llbracket 1, n \rrbracket^2$$

the $n \times n$ covariance matrix of F at evaluation points in \mathbb{S} , and

$$\mathbf{k}(\mathbf{x}) = [k(\mathbf{x}_1, \mathbf{x}), \dots, k(\mathbf{x}_n, \mathbf{x})]^\top$$

the vector of covariances between $F(\mathbf{x})$ and $\mathbf{F}_{\mathbb{S}}$

The prediction method (Matheron [1969]) assumes that the mean of $F(\mathbf{x})$ can be written as a finite linear combination

$$m(\mathbf{x}) = \boldsymbol{\beta}^\top \mathbf{p}(\mathbf{x}),$$

where $\boldsymbol{\beta}$ is a vector of fixed but unknown coefficients, and

$$\mathbf{p}(\mathbf{x}) = [p_1(\mathbf{x}), \dots, p_l(\mathbf{x})]^\top$$

is a vector of known functions of the factor vector \mathbf{x} . Usually these functions are monomials of low degree in the components of \mathbf{x} (in practice, the degree does not exceed two). These functions may be used to reflect some prior knowledge on the unknown function. As we have none for the examples considered here, we simply use an unknown constant.

The Kriging predictor at \mathbf{x} is then the best linear predictor subject to the unbiasedness constraint $\mathbb{E}(\hat{F}(\mathbf{x})) = m(\mathbf{x})$, whatever the unknown β . The unbiasedness constraint translates into

$$\beta^\top \mathbf{P}^\top \boldsymbol{\lambda}(\mathbf{x}) = \beta^\top \mathbf{p}(\mathbf{x}), \quad (15)$$

with

$$\mathbf{P} = \begin{pmatrix} \mathbf{p}(\mathbf{x}_1)^\top \\ \vdots \\ \mathbf{p}(\mathbf{x}_n)^\top \end{pmatrix}.$$

For (15) to be satisfied for all β , the Kriging coefficients must satisfy the linear constraints

$$\mathbf{P}^\top \boldsymbol{\lambda}(\mathbf{x}) = \mathbf{p}(\mathbf{x}), \quad (16)$$

called *universality constraints* by Matheron. At this point, Kriging can be reformulated as follows: find the vector of Kriging coefficients that minimizes the variance of the prediction error (14) subject to the constraints (16). This problem can be solved via a Lagrangian formulation, with $\boldsymbol{\mu}(\mathbf{x})$ a vector of l Lagrange multipliers. The coefficients $\boldsymbol{\lambda}(\mathbf{x})$ are then solutions of the linear system of equations

$$\begin{pmatrix} \mathbf{K} & \mathbf{P} \\ \mathbf{P}^\top & \mathbf{0} \end{pmatrix} \begin{pmatrix} \boldsymbol{\lambda}(\mathbf{x}) \\ \boldsymbol{\mu}(\mathbf{x}) \end{pmatrix} = \begin{pmatrix} \mathbf{k}(\mathbf{x}) \\ \mathbf{p}(\mathbf{x}) \end{pmatrix}, \quad (17)$$

with $\mathbf{0}$ a matrix of zeros. A convenient expression for the variance of the prediction error is obtained by substituting $\mathbf{k}(\mathbf{x}) - \mathbf{P}\boldsymbol{\mu}(\mathbf{x})$ for $\mathbf{K}\boldsymbol{\lambda}(\mathbf{x})$ in (14) as justified by (17), to get

$$\hat{\sigma}^2(\mathbf{x}) = \mathbb{E} \left[F(\mathbf{x}) - \hat{F}(\mathbf{x}) \right]^2 = k(\mathbf{x}, \mathbf{x}) - \boldsymbol{\lambda}(\mathbf{x})^\top \mathbf{k}(\mathbf{x}) - \mathbf{p}(\mathbf{x})^\top \boldsymbol{\mu}(\mathbf{x}). \quad (18)$$

The variance of the prediction error at \mathbf{x} can thus be computed without any evaluation of f , using (17) and (18). It provides a measure of the quality associated with the Kriging prediction. Evaluations of f remain needed to estimate the parameters of the covariance of F (if any), as will be seen in Section 8.3.2.

Once f has been evaluated at all evaluation points, the prediction of the value taken by f at \mathbf{x} becomes

$$\hat{f}(\mathbf{x}) = \boldsymbol{\lambda}(\mathbf{x})^\top \mathbf{f}_\mathbb{S}, \quad (19)$$

with $\mathbf{f}_\mathbb{S} = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]^\top$ ($\mathbf{f}_\mathbb{S}$ is viewed as a sample value of $\mathbf{F}_\mathbb{S}$).

It is easy to check that (17) implies that

$$\forall \mathbf{x}_i \in \mathbb{S}, \quad \hat{F}(\mathbf{x}_i) = F(\mathbf{x}_i).$$

The prediction of f at $\mathbf{x}_i \in \mathbb{S}$ is then $f(\mathbf{x}_i)$, so Kriging is an interpolation with the considerable advantage that it also accounts for model uncertainty through an explicit characterization of the prediction error.

Remark: The Bayesian framework (see, for instance, Williams and Rasmussen [1996]) is an alternative approach to derive the BLUP, in which the model F is viewed as a Bayesian prior on the output. In the case of a zero-mean model, the conditional distribution of the function is then Gaussian with mean

$$\mathbb{E}[F(\mathbf{x}) | \mathbf{F}_S = \mathbf{f}_S] = \mathbf{k}(\mathbf{x})^\top \mathbf{K}^{-1} \mathbf{f}_S, \quad (20)$$

and variance

$$\text{Var}[F(\mathbf{x}) | \mathbf{F}_S = \mathbf{f}_S] = k(\mathbf{x}, \mathbf{x}) - \mathbf{k}(\mathbf{x})^\top \mathbf{K}^{-1} \mathbf{k}(\mathbf{x}),$$

which are exactly the mean (19) and variance (18) of the Kriging predictor for a model F with zero mean. The Kriging predictor can also be viewed as the conditional mean of $F(\mathbf{x})$ in the case of an unknown mean, if the universality constraints are viewed as a non-informative prior on β .

8.2 Kriging when f is evaluated approximately

The Kriging predictor was previously defined as the element of the space \mathbb{H}_S generated by the random variables $F(\mathbf{x}_i)$ that minimizes the prediction error. A natural step is to extend this formulation to the case of a function whose evaluations are corrupted by additive independent and identically distributed Gaussian noise variables ε_i with zero mean and variance σ_ε^2 . The model of the observations then becomes $F_{\mathbf{x}_i}^{\text{obs}} = F(\mathbf{x}_i) + \varepsilon_i$ $i = 1, \dots, n$, and the Kriging predictor for $F(\mathbf{x})$ takes the form $\hat{F}(\mathbf{x}) = \boldsymbol{\lambda}(\mathbf{x})^\top \mathbf{F}_S^{\text{obs}}$ with $\mathbf{F}_S^{\text{obs}} = [F_{\mathbf{x}_1}^{\text{obs}}, \dots, F_{\mathbf{x}_n}^{\text{obs}}]^\top$. The unbiasedness constraint (16) remain unchanged, while the mean-square error (2) becomes

$$\mathbb{E}[\hat{F}(\mathbf{x}) - F(\mathbf{x})]^2 = k(\mathbf{x}, \mathbf{x}) + \boldsymbol{\lambda}(\mathbf{x})^\top (\mathbf{K} + \sigma_\varepsilon^2 \mathbf{I}_n) \boldsymbol{\lambda}(\mathbf{x}) - 2\boldsymbol{\lambda}(\mathbf{x})^\top \mathbf{k}(\mathbf{x}),$$

with \mathbf{I}_n the identity matrix. Finally, using Lagrange multipliers as before, it is easy to show that the coefficients $\boldsymbol{\lambda}(\mathbf{x})$ of the prediction must satisfy

$$\begin{pmatrix} \mathbf{K} + \sigma_\varepsilon^2 \mathbf{I}_n & \mathbf{P} \\ \mathbf{P}^\top & 0 \end{pmatrix} \begin{pmatrix} \boldsymbol{\lambda}(\mathbf{x}) \\ \boldsymbol{\mu}(\mathbf{x}) \end{pmatrix} = \begin{pmatrix} \mathbf{k}(\mathbf{x}) \\ p(\mathbf{x}) \end{pmatrix}. \quad (21)$$

The resulting prediction is no longer interpolative, but can still be viewed as the mean of the conditional distribution of F . The variance of the prediction error is again obtained using (18).

8.3 Covariance choice

Choosing a suitable covariance function $k(\cdot, \cdot)$ for a given f is a recurrent and fundamental question. It involves the choice of a parametrized class (or model) of

covariance, and the estimation of its parameters.

8.3.1 Covariance classes

The asymptotic theory of Kriging (Stein [1999]) stresses the importance of the behaviour of the covariance near the origin. This behaviour is indeed linked with the quadratic-mean regularity of the random process. For instance, if the covariance is continuous at the origin, then the process will be continuous in quadratic mean. In practice, one often uses covariances that are *invariant by translation* (or equivalently *stationary*), *isotropic*, and such that regularity can be adjusted. Non-stationary covariances are seldom used in practice, as they make parameter estimation particularly difficult (Chilès and Delfiner [1999]). Isotropy, however, is not required and can even be inappropriate when the factors are of different natures. An example of an anisotropic, stationary covariance class is $k(\mathbf{x}, \mathbf{y}) = k(h)$, with $h = \sqrt{(\mathbf{x} - \mathbf{y})^\top \mathbf{A}(\mathbf{x} - \mathbf{y})}$ where $(\mathbf{x}, \mathbf{y}) \in \mathbb{X}^2$ and \mathbf{A} is a symmetric positive definite matrix.

A number of covariance classes are classically used (for instance exponential $h \mapsto \sigma^2 \exp(-\theta|h|^\alpha)$, product of exponentials, or polynomial). The *Matérn covariance* class offers the possibility to adjust regularity with a single parameter (Stein [1999]). The Fourier transform of a Matérn covariance is

$$\hat{k}(\omega) = \frac{\sigma^2}{(\omega_0^2 + \omega^2)^{\nu+1/2}}, \quad \omega \in \mathbb{R}, \quad (22)$$

where ν controls the decay of $\hat{k}(\omega)$ at infinity and therefore the regularity of the covariance at the origin. Stein [1999] advocates the use of the following parametrization of the Matérn class:

$$k(h) = \frac{\sigma^2}{2^{\nu-1}\Gamma(\nu)} \left(\frac{2\nu^{1/2}h}{\rho} \right)^\nu \mathcal{K}_\nu \left(\frac{2\nu^{1/2}h}{\rho} \right), \quad (23)$$

where \mathcal{K}_ν is the modified Bessel function of the second kind (Yaglom [1986]). This parameterization is easy to interpret, as ν controls regularity, σ^2 is the variance ($k(0) = \sigma^2$), and ρ represents the *range* of the covariance, *i.e.*, the characteristic correlation distance. To stress the significance and relevance of the regularity parameter, Figure 10 shows the influence of ν on the covariance, and Figure 11 demonstrates its impact on the sample paths. Since Kriging assumes that f is a sample path of F , a careful choice of the parameters of the covariance is essential.

8.3.2 Covariance parameters

The parameters for a given covariance class can either be fixed using prior knowledge on the system, or be estimated from experimental data. In geostatistics, es-




Fig. 10. Matérn covariances for the parameterization (22) with $\rho = 0.5$, $\sigma^2 = 1$. Solid line corresponds to $\nu = 4$, dashed line to $\nu = 1$ and dotted line to $\nu = 0.25$.




Fig. 11. Three sample paths of a zero-mean Gaussian process with a Matérn covariance. Conventions are as in Figure 10: $\nu = 4$ for the solid line, $\nu = 1$ for the dashed line and $\nu = 0.25$ for the dotted line.

timization is carried out using the adequacy between the empirical and model covariances (Chilès and Delfiner [1999]). In other areas, cross validation (Wahba [1998]) and maximum likelihood (Stein [1999]) are mostly employed. For simplicity and generality reasons (Stein [1999]), the maximum-likelihood method is preferred here.

The maximum-likelihood estimate of the parameters of the covariance maximizes the probability density of the data. Using the joint probability density of the observed Gaussian vector, and assuming that the mean of $F(\mathbf{x})$ is zero for the sake of simplicity, the maximum-likelihood estimate of the vector $\boldsymbol{\theta}$ of the covariance parameters is obtained (see, for instance, Vechia [1998]) by minimizing the negative log-likelihood

$$l(\boldsymbol{\theta}) = \frac{n}{2} \log 2\pi + \frac{1}{2} \log \det \mathbf{K}(\boldsymbol{\theta}) + \frac{1}{2} \mathbf{f}_s^\top \mathbf{K}(\boldsymbol{\theta})^{-1} \mathbf{f}_s, \quad (24)$$

In the case of an unknown mean for $F(\mathbf{x})$, it is possible to estimate the parameters, using for example the *REstricted Maximum Likelihood* (REML, see Stein [1999]). This is the approach used for the examples in this paper.

Figure 12 illustrates prediction by Kriging with a Matérn covariance, the parameters of which have been estimated by REML. The prediction interpolates the data, and confidence intervals are deduced from the square root of the variance of the prediction error to assess the quality of the prediction between data. Figure 12 also contains a series of conditional simulations (obtained with the method explained in Section 3.2), namely sample paths of F that interpolate the data. As implied by (20), the Kriging prediction is the mean of these conditional simulations.

Acknowledgements

The authors wish to thank Donald R. Jones for his comments that greatly contributed to improving the accuracy and clarity of this paper.

References

- P. Abrahamsen. A review of gaussian random fields and correlation functions. Technical report, Norwegian Computing Center, 1997. URL www.math.ntnu.no/omre/TMA4250/V2007/abrahamsen2.ps.
- R.J. Adler. On excursion sets, tubes formulas and maxima of random fields. *Ann. Appl. Probab.*, 10(1):1–74, 2000.
- S. Chib and E. Greenberg. Understanding the metropolis hastings algorithm. *Amer. Statist.*, 49(4):327–335, 1995.




Fig. 12. Example of Kriging interpolation (bold line) for a function of one variable. The data are represented by squares, and the covariance parameters were estimated by REML. Dashed lines delimit 95% confidence region for the prediction. The thin solid lines are examples of conditional simulations.

- J.P. Chilès and P. Delfiner. *Geostatistics, Modeling Spatial Uncertainty*. John Wiley & Sons, Inc, New York, 1999.
- T. M. Cover and A. J. Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc, New York, 1991.
- D. Cox and S. John. Sdo: a statistical method for global optimization. In N. Alexandrov and M. Y Hussaini, editors, *Multidisciplinary Design Optimization: State of the Art*, pages 315–329, Philadelphia, 1997. SIAM. URL citeseer.ifi.unizh.ch/cox97sdo.html.
- P. Delfiner. *Shift Invariance Under Linear Models*. PhD thesis, Princetown University, New Jersey, 1977.
- L.C.W. Dixon and G.P. Szegö. The global optimisation problem: an introduction. In L.C.W. Dixon and G.P. Szegö, editors, *Towards Global Optimization 2*. North-Holland Publishing Company, 1978.
- D. Geman and B. Jedynak. An active testing model for tracking roads in satellite images. Technical Report 2757, Institut National de Recherche en Informatique et en Automatique (INRIA), December 1995.
- H.M. Gutmann. A radial basis function method for global optimization. *J. Global Optim.*, 19(3):201–227, 2001.
- D. Huang. *Experimental Planning and Sequential Kriging Optimization Using Variable Fidelity Data*. PhD thesis, Ohio State University, 2005.
- D. Huang, T. Allen, W. Notz, and N. Zeng. Global optimization of stochastic black-box systems via sequential kriging meta-models. *J. Global Optim.*, 34:441–466, 2006.

- D.R. Jones. A taxonomy of global optimization methods based on response surfaces. *J. Global Optim.*, 21:345–383, 2001.
- D.R. Jones, M. Schonlau, and J. William. Efficient global optimization of expensive black-box functions. *J. Global Optim.*, 13:455–492, 1998.
- G. Matheron. Principles of geostatistics. *Economic Geology*, 58:1246–1266, 1963.
- G. Matheron. Le krigeage universel. In *Cahiers du Centre de Morphologie Mathématique de Fontainebleau*. Ecole des Mines de Paris, 1969. Fasc. 1.
- M.J. Sasena, P. Papalambros, and P. Goovaerts. Exploration of metamodeling sampling criteria for constrained global optimization. *Eng. Opt.*, 34:263–278, 2002.
- M. Schonlau. *Computer Experiments and Global Optimization*. PhD thesis, University of Waterloo, 1997.
- E. Sjö. *Crossings and Maxima in Gaussian Fields and Seas*. PhD thesis, Lund Institute of Technology, 2000.
- A.J. Smola. *Learning with Kernels*. PhD thesis, Technische Universität Berlin, 1998.
- M.L. Stein. *Interpolation of Spatial Data: Some Theory for Kriging*. Springer, New-York, 1999.
- A.V. Vechia. Estimation and model identification for continuous spatial processes. *J. Royal Statist. Soc.*, B(50):297–312, 1998.
- G. Wahba. Support vector machines, reproducing kernel Hilbert spaces, and randomized GACV. In B. Schölkopf, C.J.C. Burges, and A.J. Smola, editors, *Advances in Kernel Methods - Support Vector Learning*, volume 6, pages 69–87, Boston, 1998. MIT Press.
- A.G. Watson and R.J. Barnes. Infill sampling criteria to locate extremes. *Math. Geol.*, 27(5):589–698, 1995.
- B.J. Williams, T.J. Santner, and W.I. Notz. Sequential design of computer experiments to minimize integrated response functions. *Statist. Sinica*, 10:1133–1152, 2000.
- C.K.I Williams and C.E. Rasmussen. Gaussian processes for regression. In D.S. Touretzky, M.C. Mayer, and M.E. Hasselmo, editors, *Advances in neural information processing systems*, volume 8. MIT Press, 1996.
- A.M. Yaglom. *Correlation Theory of Stationary and Related Random Functions I: Basic results*. Springer Series in Statistics. Springer-Verlag, New-York, 1986.